

NIST Network and Telephone Time Services

Judah Levine
Time and Frequency Division
National Institute of Standards
and Technology
Boulder, Colorado

Mailing address:

Mail Stop 847
National Institute of Standards and Technology
325 Broadway
Boulder, Colorado 80303
Voice: (303) 497 3903
FAX: (303) 497 6461
e-mail: jlevine@boulder.nist.gov
web page: www.boulder.nist.gov/timefreq
web time: www.time.gov

History of US standards

- Congress fixes Weights and Measures
 - Constitution, Article I, Section 8
- Weights and measures to Treasury
 - 14 June 1836 (5 Stat. 133)
- NBS “Organic act” 3 March 1901 (31 Stat. 1449)
- NBS to Dept. of Commerce in 1903

Section 2 of the organic act:

“That the functions of the bureau shall consist in the custody of the standards; the comparison of the standards used in scientific investigations, engineering, manufacturing, commerce, and educational institutions with the standards adopted or recognized by the Government.”

The organic act was amended in 1950 (64 Stat. 371). The amended version kept essentially the same language.

FY98 Budget Language:

Dept of Commerce/Technology Administration/NIST:

+ Encourages the development of the technological infrastructure required to support industry in the 21st century

+ Fosters the development, diffusion and adoption of new technologies

+ Disseminates technological information

+ Seeks to create a business environment conducive to innovation

The Metric Convention

- Signed in Paris on 20 May 1875
 - established BIPM in Sevres
- General Conference on Weights and Measures (CGPM) -- “Mise en pratique”
- International Committee on Weights and Measures (CIPM) -- VP from NIST
- Consultative Committee for T & F
 - working group on atomic time

The convention of the meter was ratified by President Rutherford B. Hayes on 27 September 1878 on the advice of the Senate.

The time and frequency functions were initially implemented by the Bureau International de l'heure, which was located at the Paris Observatory. These functions were transferred to the International Bureau of Weights and Measures (BIPM) in the 70's. Although this change was done at least partly for administrative reasons, it really reflected the move away from astronomical observations as the fundamental source of time. The management of leap seconds -- the only remaining connection between astronomy and precision civilian time keeping, was not moved into the BIPM and remains a separate office at the Observatory.

The BIPM Consultative Committee for Time and Frequency was originally called the Consultative Committee for the Definition of the Second. Its name was changed last year.

The CCTF and the Working Group on Atomic Time generally meet about every two years at the BIPM. They are concerned with the details of how atomic time is realized and how data are to be exchanged among the laboratories. The decisions must be ratified by the higher levels, but this is usually a formality.

International Atomic Time (TAI)

- Computed monthly by BIPM using data from a world-wide ensemble of ~250 commercial frequency standards
 - computation is about one month after the fact
- Additional data from **primary** frequency standards (NIST, PTB, ...) used to maintain long-term stability

The computation of TAI for each month is usually completed by the 15th of the following month, and sometimes a few days earlier. The results are published by the BIPM in its Circular T.

A primary frequency standard is a device that realizes the definition of the standard of frequency (which is based on a transition in cesium atoms). It operates on the same general physical principles as a commercial cesium frequency standard, but is much more carefully designed so as to ensure that its accuracy can be evaluated.

Only a few such devices exist. The current operational standard in the US is NIST-7; the PTB in Germany also has operational primary standards of roughly equivalent capability. The next generation of primary standards (based on cooled-atom technology) is now under construction at NIST and at other national timing laboratories.

NIST and many other laboratories also have longer-term research programs aimed at developing the next generation of clocks. In addition to work on improving cesium-based standards, these groups are looking at standards based on Calcium, Mercury ions and Rubidium.

Coordinated Universal Time (UTC)

- Derived from TAI by addition of leap seconds as needed (every 1-1.5 yr)
 - maintains synchronization between atomic time scales and rotation rate of the earth
 - next leap second probably in June, 2000
- Available only after the fact
- Is not realized in a physical clock
- Old “GMT” is similar but not identical

The implementation of UTC as TAI plus leap seconds dates from 1972.

Leap seconds are usually added at the end of June or December -- after 23:59:59 on 31 December, for example. The name of the leap second is 23:59:60, and the next second is 00:00:00 of the next day. The definition also includes dropping a second, but this has never been done and it is unlikely to be needed for the foreseeable future.

There is no general way for assigning a time-tag to an event that happens during a leap second -- especially for systems that keep time as the number of seconds since some epoch. The normal strategy is to effectively stop such clocks during the leap second.

No leap second is scheduled at this time, but it is likely that one will be announced next year. The difference between TAI and UTC is currently 32 s. In other words, TAI has been ahead of UTC by that amount since the last leap second at the end of December, 1998.

Since TAI (and hence UTC) represent averages based on data from many laboratories, these composite scales are not realized exactly by any clock anywhere.

UTC and GMT are essentially equivalent in a practical sense, but they are theoretically different. UTC is based on data from atomic clocks (with leap seconds added). GMT is an astronomical scale.

US estimates of UTC: UTC(NIST) and UTC(USNO)

- Derived in real-time from local clock ensemble using averaging procedure
- Steered towards UTC using small rate changes applied monthly(NIST) or irregularly as needed (USNO).
- Differences are slowly varying and are on the order of nanoseconds.
 - not always true outside of US

The NIST clock ensemble is located in Boulder, Colorado. It consists of the primary frequency standard NIST-7 and about a dozen commercial cesium standards and hydrogen masers.

The ensemble at the USNO consists of commercial devices. It has more clocks than the NIST ensemble, but has roughly comparable performance.

The steering corrections applied to realize UTC(NIST) are published in advance in the NIST Time and Frequency Bulletin. The correction is applied at 0000 UTC on the first day of each month. A typical correction is a rate change of 0.5-1 ns/day (a change of about 10^{-14} in fractional frequency). Time adjustments are never used.

1 nanosecond= 0. 000 000 000 1 s. A difference of this magnitude is not negligible for many NIST customers. In addition to guaranteeing a small time offset, NIST also tries to keep the time scale as smooth as possible. This is important for many users who depend on NIST for frequency information (NASA/DSN, telecomm ...).

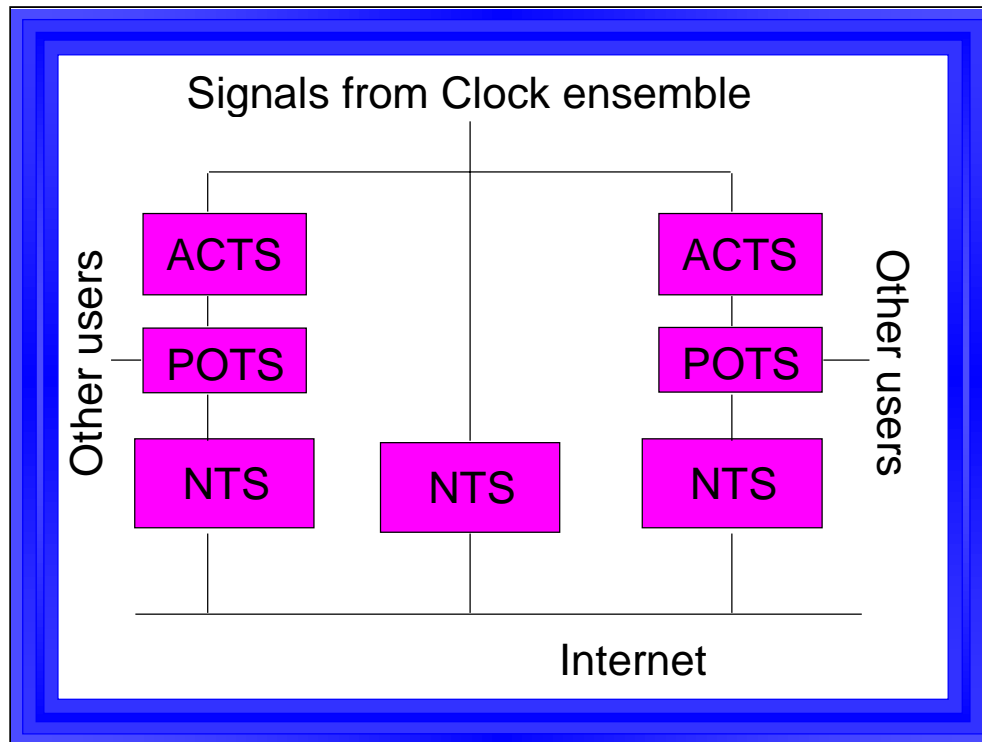
Time Services

- NIST radio services
 - WWV, WWVH, WWVB
 - Same message on dial-up “talking clock”
- ACTS
 - dial-up digital time service
- Internet services
 - NTP, daytime, time, ...
- Authenticated service (planned)

Radio station WWV is located near Fort Collins, Colorado. It transmits standard time and frequency information simultaneously on a number of short-wave frequencies. WWVH is located in Maui, Hawaii and transmits the same information on the same frequencies. Radio Station WWVB is also located near Fort Collins, Colorado. It transmits on 60 kHz and is intended for automated systems and more accurate frequency comparisons.

The ACTS service hardware is located in Boulder, Colorado. It uses standard modems and dial-up telephone lines to transmit time information with an uncertainty of a few milliseconds. NIST has written example client programs for a number of different platforms and many third-party programs are also available.

NIST currently operates 13 servers for transmitting time on the Internet in a number of different formats. In addition to the standard Network Time Protocol, NIST has developed other protocols that are better suited to PCs and small workstations and has developed client software for a number of common configurations. One of these servers transmits time messages exactly 2 years in the future. It is used for testing systems for Y2K performance.



The ACTS servers are located at the NIST laboratory in Boulder, Colorado. They are driven by a direct connection to the cesium clocks in the NIST atomic clock ensemble.

We currently operate 7 ACTS servers including a special Y2K server, which transmits time messages that are exactly 2 years in the future. The system is interfaced to the telephone network using 30 telephone lines and standard multi-speed modems.

If the user echoes the on-time marker back to NIST, the hardware at NIST measures the round trip delay and advances subsequent on-time markers to correct for it. If the user does not echo the on-time marker then a nominal fixed delay is assumed. No special hardware or software is required at the user end -- echoing the on-time marker is enough.

The Network Time servers that are located at NIST in Boulder are also directly connected to the cesium clocks in the atomic clock ensemble. The Network Time servers at other locations are synchronized using periodic dial-up connections to the ACTS system in Boulder. The NTP protocol provides a means for estimating the network delay in software -- no hardware adjustments are made.

Monitoring and Control

- ACTS and Network Monitors
 - Located in Boulder and independently synchronized to UTC(NIST)
 - Connect to each server as a “customer” and verify time message.
 - Request status logs and performance estimates from each server using private channel
 - Monitors check each other

The ACTS monitors dial each one of the servers and compare the received time message with their own internal reference clock, which is a hardware device that is driven by signals from the cesium clock ensemble.

The network time service perform the same kind of test using periodic network connections.

In addition, each server monitors its own performance, and signals an alarm condition if these internal checks fail.

The alarm system is linked to the NIST guard office and to staff pagers.


When the user echoes the on-time marker and the servers are measuring the delay, the accuracy and stability of the ACTS time messages are on the order of a few milliseconds. The network time servers are synchronized to UTC(NIST) with this uncertainty, but the messages received from the network time services are both less stable and less accurate than this because of the jitter and asymmetry in the network delay.

Typical Traceability Requirement

- National Association of Securities Dealers (NASD)
 - Order Audit Trail System (OATS)
 - Rule 6953
 - NASD Notice 98-33
 - OATS Reporting Technical Specifications, Chapter 2
 - Clocks traceable to NIST within 3s


Any of the NIST time services can provide this level of accuracy with no difficulty. The 3s tolerance can probably be realized by using a simple “set and forget” algorithm to synchronize clocks once or twice per day. More sophisticated algorithms can do much better than this -- either achieving more accurate synchronization or requiring less frequent calibrations.

The current rules do not specify how the traceability to NIST is to be verified; most clocks now keep internal log files, and this is likely to be the basis of future specifications.



The early-morning peak in the ACTS usage is caused by OATS clocks at securities dealers which are synchronized to UTC(NIST) before the markets open. Many of these devices call NIST at 4 am Eastern time, which is 0800 UTC in the Summer and 0900 UTC in the Winter. Note the one-hour shift in the most recent data because of the switch to standard time at the end of October.

In its current configuration, the ACTS system can handle about 2500 calls per hour on the average, and the peak load *averaged over a full hour* is about 40% of this value. This calculation assumes that the calls are uniformly distributed during the hour period, and this is often not the case. The ACTS system is often completely busy for one or two minutes at the start of each hour.



This shows the load on all of the NIST network time servers for the last 2.5 years. The underlying growth rate is about 7% per month compounded.

The large spikes in May of 1998 and 1999 are caused by students testing NTP client software as part of a class project.

The trace marked "other fmt" represents requests in time, daytime and icmp/ping formats. These requests make up about 15% of the total load on all of the servers. This percentage has not changed significantly since we started keeping records of the load.

Advantages of Cable-Based Time services

- Two-way communication over network can support authentication
 - simple one-way broadcasts cannot do this
- Infrastructure already present and paid for
- Good scaling properties for large networks

Authentication includes two aspects:

1. Preventing spoofing and other attacks that compromise the accuracy or the traceability of the time service messages
2. Providing a mechanism so that a user can prove to an independent third party that the time of the client system was accurate at some previous epoch. Simply “doing the right thing” may not be enough to satisfy this requirement.

Vulnerabilities of Network Time Services

- Undetected attacks are most serious
 - denial of service and reusing old messages
easy to do but easy to detect
- Spoofing of servers
 - identity of servers based on ip address
- Asymmetric Delaying of messages
 - introduces bias into delay estimator equal to one-half of asymmetry

The NTP protocol currently supports authentication using symmetric key algorithms. The sender hashes a message using one of a number of algorithms and a secret key, and the receiver authenticates the message by repeating the hashing process and comparing the local computation with the received hash value.

This method is a good first step, and is probably adequate for many private networks. It does not address all of the problems, especially for public servers such as those that are operated by NIST. It does not prevent a client for spoofing a server, for example, and the basic algorithm does not support any way for the client to prove to a third party that its clock really was synchronized at a given epoch.

A natural extension would be to use public-key methods either directly to validate the time data or indirectly to exchange a session key which is then used in the usual symmetric-key way.

Note that ACTS is more robust against both of these attacks, since it uses the closed telephone system rather than the open Internet to exchange data.

Fixing the Problems ...

- Robust authentication and maintaining an audit trail are not trivial in concept or simple to implement in practice
- Many (most) users do not need the extra cost and complexity that are implied by these requirements
- A “layered” approach may be the optimum way of realizing the solution

In addition to being a complicated problem to solve in general, different users may have very different detailed requirements with respect to the hardware that is used to implement the solution.

A layered approach would only provide full authentication and traceability to those users who needed it and were willing to pay to get it. In addition, the end-user hardware could be designed to meet the needs of many different kinds of customers.